AI SEMANTIC INFRASTRUCTURE (ASI)

A HIGH-LEVEL ARCHITECTURE FOR A GLOBAL INNOVATION INTELLIGENCE INFRASTRUCTURE

ARCHITECTURE • SEMANTICS • ARTIFICIAL INTELLIGENCE • INNOVATION

AI SEMANTIC INFRASTRUCTURE (ASI)

A HIGH-LEVEL ARCHITECTURE FOR A GLOBAL INNOVATION INTELLIGENCE INFRASTRUCTURE

First English Edition
2026

© 2026 Sirazhudin Guseynov

All rights reserved.

No part of this publication may be reproduced, distributed, or transmitted in any form or by any means without prior written permission of the author, except for brief quotations used in reviews or scholarly references.

ABOUT THIS BOOK

This book presents the high-level architecture of AI Semantic Infrastructure (ASI)—a conceptual infrastructure designed to collect, semantically process, evaluate, and route technological ideas to organizations capable of implementing them.

ASI is not a software product, nor a complete technical specification. It is an architectural proposal describing the principles, components, responsibilities, and interactions required to build a scalable global infrastructure for innovation intelligence.

The implementation of such an infrastructure may vary across organizations and deployment models, while remaining consistent with the architectural principles described in this work.

CONTENTS

Preface

Introduction

PART I. VISION AND FOUNDATIONAL PRINCIPLES

Chapter 1. The Crisis of Fragmented Technological Innovation

1.1. A World in Which Innovation Fragments
1.2. The Scalability Constraint of Human Expertise
1.3. Ideas Dissipate Prior to Realization
1.4. The Inaccessibility of Tech Giants as a Technical Constraint
1.5. The Absence of a Global Innovation Infrastructure
1.6. Why the Crisis Has Become Visible Now

Chapter 2. Philosophy and Core Principles of ASI

2.1. A New Philosophy of Innovation
2.2. The Primary Objective of ASI
2.3. The Role of AI: Processing the Inaccessible
2.4. Human-Centered Philosophy
2.5. The Priority of Innovation Principle
2.6. Semantic Evaluation as Foundation

Chapter 3. High-Level System Architecture and Objectives

3.1. ASI as a Global Innovation Infrastructure
3.2. The Central Semantic Core
3.3. Organizational Nodes
3.4. Primary System Flows
3.5. Hybrid Deployment Model
3.6. High-Level Architectural Objectives
3.7. Architectural Rationale

Chapter 4. Architectural Principles and Deliberate Constraints

4.1. Principle of the Unified Semantic Core
4.2. Principle of Autonomous Yet Compatible Organizational Nodes
4.3. Principle of Stream-Oriented Idea Processing
4.4. Principle of Human-Centered Architecture
4.5. Principle of Transparency and Explainability
4.6. Principle of Data Minimization and Privacy by Design
4.7. Principle of Scalability as an Intrinsic Property
4.8. Deliberate Constraints (Non-Goals)

PART II. IDEA INTAKE AND INITIAL PROCESSING

Chapter 5. Architecture of Idea Submission and Intake

5.1. The Idea Submission Process
5.2. Submission Lifecycle
5.3. User Interaction with the System
5.4. Idea Intake Channels

Chapter 6. Architecture of Identification, Privacy, and Anonymity

6.1. General Principles of Identification and Privacy
6.2. Fully Anonymous Submission
6.3. Authorized Submission
6.4. Authorized Anonymous Submission
6.5. Author Identity Isolation
6.6. Personal Data Protection

Chapter 7. Architecture of Pre-Evaluation and Quality Control

7.1. Purpose and Boundaries of Pre-Evaluation
7.2. Initial Idea Validation
7.3. Spam Detection and Bulk Attack Prevention
7.4. Filtering of Inappropriate and Prohibited Content
7.5. Duplicate Warning
7.6. Metadata Preparation and Anonymization

Chapter 8. Architecture of Duplicate Detection and Semantic Similarity

8.1. Definition of Semantic Duplicate
8.2. Multi-Layer Similarity Analysis
8.3. Distinction Between Purpose and Mechanism
8.4. Protection Against Copy-Paste and Obfuscation
8.5. System Behavior Upon Duplicate Detection
8.6. Feedback to the Author

PART III. SEMANTIC CORE AND INTELLIGENCE LAYER

Chapter 9. Architecture of the Central Semantic AI and Unified Evaluation Model

9.1. Definition of the Semantic Core
9.2. Unified Reasoning Model
9.3. Functions of the Central Semantic Core
9.4. Principle of Unified Evaluation

Chapter 10. Architecture of Semantic Scoring and Multi-Criteria Evaluation

10.1. Principles of Semantic Scoring
10.2. Core Evaluation Criteria
10.3. Multi-Criteria Model and Weighting
10.4. Final Score Formation Process
10.5. The Role of Scoring in the System
10.6. Balancing Global Standardization and Local Adaptation

Chapter 11. Architecture of Dynamic Novelty, Uniqueness, and Temporal Evolution

11.1. The Dynamic Nature of Novelty
11.2. Temporal Impact on Idea Evaluation
11.3. Evolution of Uniqueness
11.4. Influence of Accepted Ideas on New Evaluations
11.5. Innovation Aging Architecture

Chapter 12. Architecture of Semantic Tags, Indexing, Categorization, and Search

12.1. Automated Semantic Tag Generation
12.2. Semantic Indexing
12.3. Multi-Dimensional Categorization
12.4. Search Mechanisms
12.5. Knowledge Navigation

Chapter 13. Architecture of AI Memory, Knowledge Storage, and Evolution

13.1. Memory of Processed Ideas
13.2. Storage of Rejected Idea Extracts
13.3. Long-Term Knowledge Storage
13.4. Memory Optimization
13.5. Knowledge Base Evolution

Chapter 14. Architecture of Semantic Synthesis and Idea Amplification

14.1. Philosophy of Synthesis and Amplification
14.2. Amplification of Weak Ideas
14.3. Semantic Combination (Synthesis)
14.4. Contribution Accounting Model
14.5. Synthesis Constraints
14.6. Architectural Significance of Synthesis

PART IV. EVALUATION, REVIEW, AND HUMAN-AI INTERACTION

Chapter 15. Architecture of Evaluation and Idea Acceptance

15.1. Overall Pipeline Structure
15.2. Idea Progression Logic
15.3. Load Consideration and Dynamic Thresholds
15.4. Admission Criteria and Final Decision

Chapter 16. Architecture of Explainability and Decision Transparency

16.1. General Explainability Requirements
16.2. AI-Generated Explanations
16.3. Transparency for Different Participants
16.4. Evaluation Traceability
16.5. Explanatory Note Generation

Chapter 17. Architecture of Analyst Workflows

17.1. Role and Core Responsibilities of Analysts
17.2. Analyst Workflows
17.3. Review Queues and Prioritization
17.4. Analysis Interfaces and Analyst Workspace
17.5. Decision Support

Chapter 18. Architecture of Human-AI Interaction

18.1. Principles of Human Control
18.2. Responsibility Distribution
18.3. Collaborative Idea Evaluation
18.4. Final Decisions and Oversight

Chapter 19. Architecture of Feedback and User Experience

19.1. General Principles of Feedback
19.2. User Notifications
19.3. Review Results
19.4. Rating Information
19.5. Author Communication

PART V. TRUST, GOVERNANCE, AND RESILIENCE

Chapter 20. Architecture of the Reputation System

20.1. Reputation System Principles
20.2. Reputation Award Mechanics
20.3. Positive Scores and Their Meaning
20.4. Negative Scores
20.5. Synthesis Contribution Accounting
20.6. Rating Structure and Contribution Quality

Chapter 21. Architecture of Abuse Protection and Adversarial Defense

21.1. General Security Principles
21.2. Reputation Inflation Protection
21.3. Coordinated Attack Countermeasures
21.4. Manipulation Detection
21.5. Handling Malicious Submissions
21.6. System Integrity Protection

Chapter 22. Architecture of Governance, Organizational Flexibility, and Corporate Configuration

22.1. Core Governance Principles
22.2. Organizational Policies
22.3. Configurable Evaluation Parameters
22.4. Governance Models
22.5. Operational Flexibility

Chapter 23. Architecture of Audit, Transparency, Security, and Access Control

23.1. Audit Logs
23.2. Transparency Mechanisms
23.3. Security Principles
23.4. Access Control
23.5. Accountability and Regulatory Compliance

PART VI. OPERATION, SCALING, AND SYSTEM EVOLUTION

Chapter 24. Architecture of Queues and High-Load Infrastructure

24.1. Processing Large Idea Streams
24.2. Queue Management
24.3. Resource Allocation and Dynamic Scaling
24.4. Computational Efficiency
24.5. Infrastructure Scalability

Chapter 25. Architecture of ASI as a Global Semantic Infrastructure

25.1. ASI as a Global Infrastructure Architecture
25.2. Organization Onboarding
25.3. Idea Submission Modes
25.4. Individual Interaction with ASI
25.5. ASI Node Autonomy
25.6. Global Nature of the Infrastructure
25.7. Freedom of Operator Choice
25.8. Boundary of the ASI Architecture

Chapter 26. Architecture of Data Storage and Information Lifecycle Management

26.1. Storage Policies
26.2. Data Retention Periods
26.3. Archiving and Deletion
26.4. Information Lifecycle Management
26.5. System Semantic Memory Management

Chapter 27. Architecture of Long-Term Learning and System Evolution

27.1. Knowledge Accumulation
27.2. Intellectual Capability Development
27.3. Semantic Model Adaptation
27.4. Long-Term Optimization
27.5. Growth of Innovation Intelligence

Chapter 28. Architecture of the Innovation Lifecycle and Semantic Knowledge Base

28.1. Idea Lifecycle
28.2. Knowledge Accumulation
28.3. Innovation Tracking
28.4. Semantic Knowledge Base Formation
28.5. Organizational Expertise Development

PART VII. ASI ECOSYSTEM AND OPERATING PRINCIPLES

Chapter 29. Author Participation, Operators, and ASI Operating Principles

29.1. Principle of Voluntary Participation
29.2. Author Participation Models
29.3. Value for Authors
29.4. Role of Operators
29.5. Diversity of Operating Models
29.6. Federative Future
29.7. Strategic Purpose of ASI

Conclusion. The Future of Global Innovation Intelligence

30.1. The Problem ASI Solves
30.2. Innovation as a Global Resource
30.3. Democratization of Innovation
30.4. From Information Networks to Innovation Networks
30.5. The Role of AI in Collective Intelligence
30.6. Long-Term Vision
30.7. Final Word

Afterword

Appendices

Appendix A. Idea Evaluation Criteria System
Appendix B. Idea Evaluation Examples
Appendix C. Examples of Semantic Synthesis
Appendix D. Reputation System Mechanics
Appendix E. Semantic Tags, Indexing, and Search
Appendix F. ASI Use Cases
Appendix G. Typical ASI Configurations
Appendix H. Comparison with Existing Systems
Appendix I. ASI Risks and Limitations

Preface

Large-scale systemic failures often originate from seemingly trivial omissions.

Several years ago, an attempt was made to communicate a technical proposal. The proposal itself was not a breakthrough—it was an incremental improvement to a routine operational system. It seemed reasonable to assume that a straightforward mechanism existed for conveying such proposals to the relevant decision-makers.

No such mechanism existed.

There was no channel, no formalized process, no institutional infrastructure through which an individual contributor could be heard. The proposal was never captured, evaluated, or routed.

Initially, this appeared to be an isolated anomaly. Further examination revealed a consistent pattern: this was not an exception but rather the default behavior of the system.

Humanity produces a substantial volume of technological ideas on a daily basis. However, the overwhelming majority of these ideas never reach the entities possessing the resources, domain expertise, and execution capacity required for implementation.

The fundamental constraint is not a scarcity of ideas.
The fundamental constraint is the absence of an infrastructure designed to systematically identify, assess, preserve, and channel those ideas toward viable realization pathways.

Global information transmission networks have been established.
Advanced knowledge storage and retrieval systems have been deployed.
However, no analogous infrastructure exists for processing the innovation signal.

This work proposes an architectural framework for such an infrastructure.

AI Semantic Infrastructure (ASI) is neither a crowdsourcing platform nor a social network. It constitutes a systemic architectural proposal: a semantic middleware layer positioned between distributed human creative output and the institutional execution capacity required for implementation.

This work does not present itself as a definitive solution. It outlines a potential infrastructure architecture—one in which the valuation of an idea is decoupled from

the author's social status, professional network, or rhetorical volume, and is instead determined by the idea's intrinsic properties: its functional logic, operational mechanics, and projected impact.

Should such an infrastructure be instantiated, any individual originating a viable idea would be positioned to state:

"My proposal has a viable pathway to consideration."

Introduction

The contemporary technological landscape produces ideas at an unprecedented velocity. Daily, independent engineers, researchers, developers, and enthusiasts generate thousands of concepts—ranging from incremental optimizations to potentially transformative breakthroughs. The overwhelming majority of these ideas, however, never reach the entities capable of evaluating, developing, or implementing them.

This is not a creative problem. It is a structural one.

Robust tools exist for information dissemination, knowledge storage, and data retrieval. Yet a specialized infrastructure—one capable of systematically collecting, semantically interpreting, evaluating, and routing technological ideas toward appropriate implementation points—remains conspicuously absent.

Existing solutions address only isolated facets of this challenge:

- Corporate idea banks — restricted to internal participation
- Crowdsourcing platforms — limited by superficial engagement metrics
- Patent systems — constrained by formal legal criteria
- Scientific publications — slow, siloed, and ill-suited for unstructured idea flow

Consequently, valuable innovation signals are lost in noise, and the potential of distributed human intelligence remains only fractionally realized.

This book presents AI Semantic Infrastructure (ASI)—an architectural proposal for a new infrastructure designed to process the global flow of technological ideas.

ASI is not a social network, not a voting platform, and not a corporate innovation management tool. It is a semantic middleware layer that:

- Accepts ideas from any individual, in any format
- Transforms them into structured semantic objects
- Evaluates novelty, mechanistic uniqueness, and practical utility
- Identifies opportunities for synthesis and amplification
- Routes the most promising solutions to organizations and teams capable of execution

The book systematically describes the architecture of such a system—from high-level vision and foundational principles to the detailed design of individual components:

- Idea submission and ingestion
- Multi-stage evaluation pipelines
- Semantic memory and knowledge persistence
- Synthesis and recombination engines
- Reputation and trust mechanisms
- Abuse protection and adversarial resilience
- Deployment models and institutional integration

This work does not present ASI as a turnkey solution. It offers an architectural approach—a foundation upon which global or corporate innovation intelligence infrastructures may be constructed. The book addresses both the strengths of the proposed architecture and its deliberate limitations, associated risks, and open questions.

Ultimately, ASI represents an attempt to answer a simple yet fundamental question:

How can humanity's best technological ideas avoid dissipation and instead find viable pathways to implementation?

PART I. Vision and Foundational Principles

Chapter 1. The Crisis of Fragmented Technological Innovation

1.1. A World in Which Innovation Fragments

Technological progress has never been faster.

It has also never been more fragmented.

Today, every corporation, laboratory, research group, and startup operates within its own cognitive silo. They do not observe what occurs beyond their boundaries—not because they lack willingness, but because the connective infrastructure between innovation nodes does not exist.

Ideas emerge everywhere.

They do not coalesce into a unified system.

They do not amplify one another.

They do not constitute a collective innovation intelligence.

Humanity generates enormous innovative potential.

Only a fraction is utilized.

This is not a failure of talent.

This is a failure of architecture.

1.2. The Scalability Constraint of Human Expertise

Prior to AI, a fundamental and uncontested boundary existed:

Expertise does not scale.

No organization, regardless of size, is capable of:

- Processing millions of ideas
- Evaluating them objectively and consistently
- Comparing them against one another at global scale
- Identifying latent semantic relationships
- Distinguishing genuine novelty from the stream of the already-known

This is not a human failing.

It is a physical constraint of biological cognition and organizational structure.

Human knowledge doubles at an accelerating rate.

Human cognitive throughput remains constant.

We have already crossed the threshold beyond which no individual—however gifted—can maintain a coherent view of the entire technological landscape.

AI does not replace the expert.

AI scales expertise to a level inaccessible to biological substrates.

1.3. Ideas Dissipate Prior to Realization

Daily, thousands of ideas capable of transforming industries, markets, technologies, and lives are generated worldwide.

Most of them dissipate.

Not because they lack merit.

But because:

- There is nowhere to submit them
- No one to evaluate them
- Nothing to compare them against
- No one to develop them further

An idea without a receptive infrastructure is a radio signal without an antenna.

It exists. No one receives it.

This book addresses the construction of an antenna for innovation.

1.4. The Inaccessibility of Tech Giants as a Technical Constraint

A persistent myth suggests: "Large corporations are closed to external ideas."

This is not social injustice.

This is a technical impossibility of the pre-AI era.

No company, regardless of scale, could:

- Process millions of incoming submissions
- Comprehend their semantic meaning
- Relate them to active projects and roadmaps
- Objectively identify valuable signals amid noise

Tech giants are not "closed."

They were physically incapable of hearing the entire world.

AI removes this constraint for the first time.

Infrastructure, however, has not followed.

1.5. The Absence of a Global Innovation Infrastructure

The core problem is not people.

Not corporations.

Not motivation.

The core problem is the absence of infrastructure.

We have internet for information.

HTTP for pages.

SMTP for messages.

DNS for domains.

We have no semantic protocol for innovation.

No system capable of:

- Accepting ideas from any individual
- Evaluating them semantically, not formally
- Comparing them against the global body of knowledge
- Amplifying weak ideas through combinatorial synthesis
- Clustering semantically related concepts
- Routing the strongest candidates to organizations capable of execution

ASI (AI Semantic Infrastructure) is an attempt to design such a protocol.

Not merely an idea repository.

But an infrastructure for innovation intelligence.

1.6. Why the Crisis Has Become Visible Now

The emergence of AI has made feasible what, five years ago, remained technical fiction:

- Real-time processing of massive idea streams
- Semantic comprehension beyond keyword matching
- Scaling expertise to global levels
- Objective, reproducible evaluation
- Synthesis and recombination based on meaning, not form

For the first time in human history, the innovation crisis has ceased to be a philosophical problem.

It has become a technical problem.

Solvable.

Architecturally.

Systemically.

This book is a project for its resolution.

Chapter 2. Philosophy and Core Principles of ASI

2.1. A New Philosophy of Innovation

The crisis of fragmented innovation cannot be resolved through methods of the past.

No single organization.

No single state.

No single expert group.

The volume of ideas generated daily worldwide exceeds the cognitive throughput of all existing institutions combined.

ASI rests on a different philosophy.

The Old Paradigm:

- Innovation is a rare event
- Talent is concentrated within corporations
- Expertise is bounded by the human brain
- Evaluation is formal—dependent on text length, structure, author status
- Infrastructure is local

The ASI Paradigm:

- Innovation is a continuous flow
- Talent is globally distributed
- Expertise scales through AI
- Evaluation is semantic—based on meaning, novelty, and value

- Infrastructure is global

This philosophy is not a declaration.

It defines every layer of the architecture.

2.2. The Primary Objective of ASI

The system has two fundamental objectives.

First: processing massive idea streams.

What was previously technically impossible: millions of ideas per day, across dozens of languages, spanning multiple industries, varying in quality and novelty.

AI makes this feasible. Not in theory. In architecture.

Second: talent discovery through ideas.

An idea is not merely text. An idea is a trace of intelligence.

When an individual consistently generates strong, deep, unconventional ideas—this is not chance. It is a cognitive pattern.

ASI transforms ideas into signals about the author's thinking, into a profile of their cognitive strengths, into a measurable indicator of potential.

Organizations gain not only ideas.

They gain the individuals capable of generating them.

2.3. The Role of AI: Processing the Inaccessible

At the core of ASI's philosophy lies a simple and rigorous premise.

AI is not needed to replace experts.

AI is needed to process what humans cannot process.

First function: scaling processing capacity.

AI can read millions of ideas, comprehend their semantic meaning, compare them globally, identify non-obvious connections, and detect novelty against the entire knowledge corpus.

This is unattainable without AI.

Second function: semantic evaluation.

AI evaluates not form. It evaluates meaning: the idea's objective, its mechanism of implementation, its novelty relative to known solutions, its potential value, its risks and limitations, and its applicability across domains.

This renders evaluation objective and scalable.

Third function: talent identification through patterns.

AI captures idea depth, temporal consistency of quality, originality capacity, systemic thinking capability, and cross-domain synthesis ability.

ASI becomes more than an idea filter.

It becomes a global talent radar.

2.4. Human-Centered Philosophy

Despite the scale of AI, ASI is built around the human.

The system comprises four roles. These are not hierarchical—they represent a distribution of responsibility.

Role	Function
Human Author	Source of ideas
AI	Scalable processing and semantic evaluation
Analyst	Final interpreter and contextual adjudicator
Organization	Executor of selected innovations

The system does not replace humans.

The system amplifies them.

2.5. The Priority of Innovation Principle

ASI proceeds from one hard rule:

If an idea has value—it must be noticed. Irrespective of its author.

From this, architectural requirements follow:

- No entry barriers
- No privileges based on author status
- No "correct" or "incorrect" author categories
- Equal infrastructure access for any individual

The priority of innovation principle is not an ethical position.

It is an engineering precondition for globality.

A system with privileges cannot be global.

2.6. Semantic Evaluation as Foundation

ASI employs no formal criteria.

No text length validation.

No description structure requirements.

No keyword matching.

No consideration of author education level.

No institutional affiliation checks.

No weight given to position or credentials.

Only meaning.

Semantic evaluation encompasses:

- Understanding the idea's objective
- Understanding its mechanism of implementation
- Understanding novelty relative to the known corpus
- Understanding the context of origin and applicability
- Understanding potential value and associated risks

This is the foundation of the entire architecture.

Everything built above depends on this layer.

Chapter 3. High-Level System Architecture and Objectives

ASI was designed not as an application or platform, but as infrastructure—a foundational layer upon which diverse innovation processes may be constructed.

3.1. ASI as a Global Innovation Infrastructure

ASI is not a corporate idea management solution, nor is it another idea-sharing platform. It is an infrastructural layer—comparable in significance to the internet or global data exchange systems.

Its purpose is to create a space in which technological ideas can flow freely from any individual, undergo deep semantic processing, and be routed to the organizations and teams capable of developing and implementing them. The system must perform equivalently well in both global open mode and fully closed corporate environments.

Deployment model flexibility is embedded in the architecture from the outset, not added retroactively. This is a fundamental requirement.

3.2. The Central Semantic Core

At the foundation of ASI lies a unified semantic core—the centralized innovation intelligence of the system.

The semantic core is not a specific AI model, but a unified system of semantic processing and evaluation that may be implemented using various artificial intelligence models. The primary requirement is strict adherence to unified rules for understanding, evaluating, comparing, and synthesizing ideas. This ensures consistency and objectivity regardless of which models are employed in the implementation.

Architecturally, ASI is not tied to any single AI provider. The semantic core defines interfaces, standards, and contracts, enabling different models to be used for different tasks—for example, a lighter model for pre-evaluation and a more powerful one for deep analysis—while maintaining uniform evaluation quality standards across the system.

The core is responsible for five key functions:

1. Deep semantic understanding of ideas — analysis of meaning, purpose, implementation mechanism, and context, not merely keywords or text structure
2. Multi-criteria evaluation and scoring — objective determination of novelty, feasibility, potential value, and risks

3. Similarity and duplicate detection — independent of language, phrasing, or description style
4. Semantic indexing and tagging — automatic classification and organization of ideas into a coherent space
5. Long-term innovation memory — accumulation and evolution of knowledge about accepted and synthesized ideas (within permitted policies)

It is through the unified semantic core that the system ensures consistency of evaluation and interpretation of ideas regardless of source or underlying AI models.

3.3. Organizational Nodes

Surrounding the semantic core are organizational nodes—connection points through which companies and research groups interface with the ASI infrastructure.

Each node may:

- Accept ideas addressed specifically to that organization
- Apply its own filtering, prioritization, and integration rules
- Interact with internal systems (R&D platforms, task trackers, product roadmaps)
- Manage data exchange levels with the global core

Nodes need not be identical. The architecture supports a wide spectrum of configurations—from fully isolated private installations to deeply integrated nodes of the global network.

3.4. Primary System Flows

The ASI architecture is organized around two main interconnected flows.

Idea Flow:
Idea enters the system → undergoes preprocessing and filtering → passes to the semantic core for deep analysis → receives scoring, tags, relationships to other ideas, and routing recommendations → enters queues of relevant organizational nodes.

Decision Flow:
Organization (via analysts or automated rules) makes a decision on the idea → result is returned to the author → accepted and synthesized ideas update the system's semantic memory. The system thus becomes self-learning: evaluation and routing quality improves over time.

3.5. Hybrid Deployment Model

The ASI architecture is designed as hybrid from inception and supports three primary modes:

Mode	Description
Global Network	Unified semantic core with multiple connected organizational nodes — the "internet of innovation"
Isolated Installations	Fully autonomous deployments within a single organization or jurisdiction
Hybrid Scenarios	Partial data exchange (with anonymization and control) between local nodes and the global core

The same architecture supports all three modes without fundamental modification.

3.6. High-Level Architectural Objectives

ASI addresses the following strategic objectives:

1. Aggregate and process the global flow of technological ideas within a unified infrastructure
2. Provide scalable semantic evaluation inaccessible to human cognition alone
3. Create an efficient mechanism for targeted delivery of ideas to relevant organizations
4. Identify talented authors through persistent patterns in their ideas
5. Preserve maximum participant flexibility and sovereignty — each organization independently determines its integration depth

3.7. Architectural Rationale

This specific architecture—with a unified semantic core, flexible organizational nodes, and support for multiple deployment models—enables resolution of the innovation fragmentation problem without disrupting existing corporate processes, instead augmenting them with a powerful semantic layer.

This is not a theoretical construct. It is a practical, implementable architecture that can be built today.

Chapter 4. Architectural Principles and Deliberate Constraints

Good architecture is defined not only by what a system can do, but also by what it consciously does not do. ASI is a large-scale infrastructure operating on massive streams of ideas and data. To preserve stability, predictability, and trust, it is built upon clear principles and hard constraints.

4.1. Principle of the Unified Semantic Core

At the center of ASI resides a unified semantic core — the system's common intelligence, ensuring consistent interpretation of ideas regardless of source, language, or domain.

This core is responsible for:

- Deep semantic understanding of idea meaning
- Unified approach to evaluation, scoring, and novelty determination
- Automated tagging, indexing, and search
- Maintenance of long-term innovation memory

The unified core guarantees that the same idea will be interpreted identically at any point within the system. This is a foundational decision that eliminates fragmentation of understanding and creates a shared innovation context.

4.2. Principle of Autonomous Yet Compatible Organizational Nodes

Organizations connect to ASI through organizational nodes—logical or physical integration points.

Each node remains autonomous: the organization independently determines policies, filtering rules, integration depth with internal systems, and data exchange levels with the global core. All nodes simultaneously maintain compatibility with semantic core standards.

This approach combines corporate sovereignty and security with the ability to participate in a global innovation network.

4.3. Principle of Stream-Oriented Idea Processing

ASI is architected as a pipeline—a sequential and scalable processing workflow for each idea:

1. Submission and initial validation
2. Pre-evaluation
3. Deep semantic analysis
4. Duplicate and relationship detection
5. Routing and scoring
6. Analytical review
7. Feedback and system memory update

The stream-oriented architecture ensures predictability, parallel processing, and resilience under high load—the only viable approach for operating at the scale of millions of ideas.

4.4. Principle of Human-Centered Architecture

Despite the central role of AI, ASI remains a human-centered system. Roles are clearly distributed:

Role	Function
Author	Idea generator
AI	Scalable semantic processing
Analyst	Contextual adjudicator and final interpreter
Organization	Executor of selected innovations

The system amplifies human capabilities but does not replace human judgment where it is critically required.

4.5. Principle of Transparency and Explainability

Every significant system decision—evaluation, routing, synthesis—must be explainable and traceable. Authors, analysts, and organizations have access to AI decision justifications. Transparency is not an add-on feature but a fundamental architectural requirement essential for trust.

4.6. Principle of Data Minimization and Privacy by Design

ASI adheres to the principle of least necessary data. The system collects and retains only information genuinely required for semantic evaluation and operation. Full anonymity is supported at the architectural level, not as an option. Privacy is embedded in the design, not layered on top.

4.7. Principle of Scalability as an Intrinsic Property

Scalability is not a goal to be achieved later, but an inherent property of the architecture. The system was designed from the outset for processing millions of ideas, thousands of organizations, and global-scale loads. All key components are therefore distributed, asynchronous, and modular.

4.8. Deliberate Constraints (Non-Goals)

To preserve stability and trust, ASI incorporates the following deliberate constraints:

- The system does not make decisions on behalf of organizations. It provides evaluation and recommendations, but final decisions always remain with humans or organizations.
- The system does not interfere with internal corporate processes. Organizational nodes are fully autonomous.
- The system does not replace experts. It scales their capabilities but does not substitute human judgment.
- The system does not impose a unified governance model. Each organization selects its preferred operational mode.
- The system does not collect or utilize excess data. Only what is necessary to perform its function.

These principles and constraints are not philosophical aspirations—they are engineering decisions that protect the system from chaos, abuse, and loss of trust. They define not only what ASI is, but also what it will never become.

PART II. Idea Intake and Initial Processing

Chapter 5. Architecture of Idea Submission and Intake

The idea submission module serves as the entry gateway to the entire AI Semantic Infrastructure. It is here that the system forms its first impression on the user and establishes the foundation for quality throughout all subsequent processing.

5.1. The Idea Submission Process

The submission process is designed to be maximally simple and accessible, with minimal barriers for the author:

1. Initiation — The user opens the submission interface via web client, mobile application, or corporate portal.
2. Idea Entry — The idea is described in free-form text. The system does not require mandatory templates or rigid structure.
3. Pre-Evaluation — Before submission, a local pre-evaluation (see Chapter 7) is triggered, analyzing text quality, identifying obvious issues, and generating initial metadata.
4. Confirmation and Submission — After reviewing pre-evaluation results, the user decides whether to proceed with submission.
5. Submission Recording — The idea receives a unique identifier and is passed to the initial processing system.

5.2. Submission Lifecycle

Each submitted idea undergoes a standardized lifecycle:

Stage	Description
Creation	Assignment of a unique ID and recording of initial state
Pre-Processing	Local quality validation, metadata generation, and initial tagging
Transfer to Semantic Core	Upon passing initial filters, the idea proceeds to deep evaluation
Status Assignment	During processing, the submission receives one of the following statuses: *In Processing*, *Transferred to Analysts*, *Accepted*, *Rejected*, *Included in Synthesis*, among others

Completion	Upon final decision, the author receives notification, and the idea is either retained in the relevant section of system memory or archived/deleted according to policy

5.3. User Interaction with the System

The idea submission interface is intentionally minimalist and oriented toward author convenience:

- Free-form text field with no rigid structural requirements
- Button to trigger pre-evaluation
- Clear system warnings and recommendations
- Submit button, activated only after successful pre-evaluation
- Personal dashboard with submission history and current statuses
- Notification system for milestone updates and final results

This approach reduces cognitive load on the user and improves the quality of incoming ideas.

5.4. Idea Intake Channels

For maximum accessibility and integration, ASI supports multiple submission channels:

Channel	Purpose
Public Web Interface	Primary channel for independent authors
Mobile Application	Simplified submission from any location
Corporate Portals	Integration within organizational environments
APIs	Automated submission from external systems and services

Specialized Integrations	Connection to internal chatbots, R&D platforms, and project management systems

This diversity of channels enables the system to collect ideas from a broad range of sources while maintaining unified processing standards.

Chapter 6. Architecture of Identification, Privacy, and Anonymity

One of the fundamental design decisions in ASI is the strict separation of author identity from the semantic content of ideas. This architectural boundary ensures a high level of privacy, reduces the risk of abuse, and strengthens trust in the system.

6.1. General Principles of Identification and Privacy

The system is designed from the ground up to maximize the separation between author identity and idea text. Identification is used only where strictly necessary—for reputation management, submission limits, and compliance with corporate policies.

Key principles:

- Data minimization — The system collects only the minimum required personal information.
- Circuit isolation — Identity, reputation, and semantic data are processed in separate, strictly partitioned modules.
- Anonymity by default — Available in certain operational modes.
- Non-recoverability — Identity cannot be reconstructed from semantic data or evaluations.

6.2. Fully Anonymous Submission

In fully anonymous mode, the system creates and stores no author identifiers whatsoever.

- User does not undergo authentication
- Idea receives only a unique submission ID
- Pre-evaluation and central AI process the idea text exclusively
- Analysts see only idea content, with no author information attached

- Author receives no feedback and has no access to history of such submissions

This mode provides maximum identity protection but excludes the possibility of reputation accumulation and result tracking.

6.3. Authorized Submission

Authorized submission is the primary mode for users seeking long-term participation in the system.

- User logs into their account
- Idea is linked to their reputation profile
- System can send notifications, account for contributions in ratings, and display submission history

Author identity remains isolated from the semantic core and analysts. The central AI and analysts work only with reputation identifiers, not with actual personal data.

6.4. Authorized Anonymous Submission

This hybrid mode allows an authorized user to submit an idea without linking it to their profile:

- Idea does not affect author reputation
- Does not appear in their personal history
- Central AI and analysts receive it as fully anonymous

This mode is useful when an author wishes to maintain privacy for a specific idea while retaining access to other system functions.

6.5. Author Identity Isolation

ASI implements multi-layer isolation:

Circuit	Function
Identity Circuit	Stores personal data (login, email, etc.)
Reputation Circuit	Operates with anonymized reputation profiles

Semantic Circuit	Processes only idea texts and their semantic objects
Analyst Circuit	Receives ideas without identity linkage

Strict technical boundaries exist between circuits. Even during authorized submission, the semantic core and analysts have no access to author personal data.

6.6. Personal Data Protection

The system adheres to the principle of least necessary data:

- Personal information never enters semantic memory and is never used in idea evaluation
- The author-idea link exists only within the identity and reputation circuits
- All data is stored in accordance with organizational policies and applicable regulatory requirements
- Full deletion of personal data upon user request is supported

This architecture ensures a high level of privacy while preserving the capacity for long-term contribution tracking and reputation management.

Chapter 7. Architecture of Pre-Evaluation and Quality Control

Pre-evaluation is the first intelligent filter of ASI, executed by an isolated lightweight AI module. Its primary purpose is to protect the system from low-quality and malicious traffic, minimize load on the central semantic core, and provide the author with rapid feedback.

7.1. Purpose and Boundaries of Pre-Evaluation

The pre-evaluation module addresses the following objectives:

- Protection of the central AI from noise, spam, and prohibited content
- Early warning to the author regarding issues with the idea
- Preparation of minimal metadata for subsequent processing
- Ensuring correct anonymization based on the selected submission mode

Important architectural constraint: The pre-evaluation module is isolated from the central semantic core and has no access to global memory. It does not influence novelty assessment, clustering, synthesis, or final ranking.

7.2. Initial Idea Validation

At the first stage, pre-evaluation analyzes the basic characteristics of the idea:

- Presence of minimal semantic structure
- Text correctness and readability
- Relevance to the domain of technological innovation
- Absence of obvious technical formatting errors

Ideas failing this validation are blocked at the entry point without transmission to the central core.

7.3. Spam Detection and Bulk Attack Prevention

Pre-evaluation identifies typical spam indicators:

- Repetitive patterns and auto-generated text
- Excessively high submission frequency from a single source
- Technical signatures (IP, device fingerprint, etc.)
- Indicators of coordinated activity

Upon spam detection, the idea is blocked. In authorized mode, the author receives a warning and may be subject to temporary restrictions.

7.4. Filtering of Inappropriate and Prohibited Content

The system automatically blocks:

- Texts containing profanity, threats, or harassment
- Malicious instructions and potentially dangerous content
- Political trolling and extremist materials
- Completely meaningless or irrelevant messages

Filtering is performed through a combination of rules and a lightweight AI model optimized for maximum speed and minimal resource consumption.

7.5. Duplicate Warning

Pre-evaluation performs a surface-level analysis for potential duplication:

- Structural comparison and key text fragments

- Basic semantic similarity

When a potential duplicate is detected, the system warns the user and suggests refinement or reformulation. The final duplicate determination is made by the central AI (see Chapter 8).

7.6. Metadata Preparation and Anonymization

Depending on the submission mode, the pre-evaluation module prepares and transmits different data sets:

Submission Mode	Data Transmitted
Fully Anonymous	Idea text only + Unique idea ID
Authorized Anonymous	Idea text only + Unique idea ID
Authorized Open	Idea text + Unique idea ID + Anonymized reputation ID + Additional metadata (per organizational policy)

Personal data (name, email, etc.) is never transmitted to the semantic core or the analyst circuit.

Chapter 8. Architecture of Duplicate Detection and Semantic Similarity

Duplicate detection is one of the key mechanisms for maintaining the quality and integrity of ASI's semantic memory. Without effective duplicate management, the system would rapidly fill with repeated solutions, diminishing the value of the global innovation memory and increasing the load on analysts.

8.1. Definition of Semantic Duplicate

In ASI, a duplicate is defined as an idea that proposes the same implementation mechanism for solving a problem, even if the phrasing, style, or specific details differ.

The primary criterion is mechanism overlap, not purpose or topic. Two ideas may address the same problem through fundamentally different approaches—in such cases, they are considered unique.

8.2. Multi-Layer Similarity Analysis

The system employs a three-layer duplicate detection model:

Layer	Scope	Method
Layer 1: Surface Analysis	Fast filter	Keyword comparison, text structure analysis, and low-level embeddings. Applied during pre-evaluation to filter obvious repetitions.
Layer 2: Deep Semantic Analysis	Central AI evaluation	Assessment of idea meaning, context, objectives, and constraints. Determines semantic proximity independent of phrasing.
Layer 3: Mechanism Analysis (Decisive)	Highest priority	Comparison of: architectural principles of the solution; sequence of steps and processes; technical and operational mechanisms; assumed pathways to achieving results.

If the implementation mechanisms substantially overlap—the idea is classified as a duplicate.

8.3. Distinction Between Purpose and Mechanism

The system clearly distinguishes three cases:

Case	Classification
Same purpose, different mechanisms	Ideas are considered unique—both may be accepted
Different purposes, same mechanism	Ideas are classified as duplicates—mechanism takes precedence over purpose
Same purpose, same mechanism	Classic duplicate

This approach preserves diversity of architectural solutions while protecting against repetition of already-known mechanisms.

8.4. Protection Against Copy-Paste and Obfuscation

Pre-evaluation and the central AI jointly protect the system against:

- Direct text copying
- Paraphrased copy-paste
- Auto-generated variations of the same idea
- Mass attempts to pollute memory

8.5. System Behavior Upon Duplicate Detection

When an idea is classified as a duplicate:

- It is rejected at the central evaluation stage
- It is not retained in long-term memory
- It is not transmitted to analysts
- It does not affect the author's reputation

Exception: If the new idea proposes a meaningful improvement to the mechanism—cost reduction, efficiency gain, implementation simplification, etc.—it is treated as an independent submission and may be accepted.

8.6. Feedback to the Author

- Authorized submission: The author receives a notification indicating the reason for rejection and, where applicable, a reference to a related idea (in anonymized form).
- Anonymous modes: No feedback is provided.

PART III. Semantic Core and Intelligence Layer

Chapter 9. Architecture of the Central Semantic AI and Unified Evaluation Model

The central semantic AI is the core intellectual component of ASI, responsible for deep understanding of ideas, their evaluation, novelty detection, and synthesis support.

9.1. Definition of the Semantic Core

The central semantic AI—or central core—is not a specific language model, but rather an architectural protocol and a set of interfaces that define unified rules for semantic processing of ideas.

This core may be implemented using different AI models (existing or future). The critical requirement is strict adherence to a unified reasoning protocol and compatibility with the remaining system components. This approach ensures vendor independence and enables system evolution over time.

9.2. Unified Reasoning Model

The core employs a unified reasoning model applied to all ideas without exception—regardless of organization, author, domain, or language.

Key properties of the unified model:

Property	Description
Universality	Operates on ideas from any domain
Formulation Independence	Processes meaning, not text

Style and Language Invariance	Independent of phrasing and language
Logical Determinism	Identical ideas yield comparable results
Transparency	Every decision is explainable

9.3. Functions of the Central Semantic Core

The central AI performs the following key functions:

- Semantic validation of the idea
- Deep analysis of the implementation mechanism
- Novelty and uniqueness assessment
- Multi-criteria scoring
- Identification of synthesis and amplification opportunities
- Preparation of explanations for analysts
- Update of system semantic memory

9.4. Principle of Unified Evaluation

All ideas undergo the same evaluation process under unified rules. This is a fundamental principle of ASI, ensuring:

- Absence of bias
- Comparability across ideas
- Fairness and system trustworthiness
- Scalability

Unified evaluation encompasses analysis of meaning, implementation mechanism, novelty, utility, feasibility, and potential.

Chapter 10. Architecture of Semantic Scoring and Multi-Criteria Evaluation

Semantic scoring is the mechanism by which the central semantic AI transforms holistic understanding of an idea into a structured, quantitatively measurable evaluation. It ensures comparability of ideas, objectivity of decisions, and transparency of the process.

10.1. Principles of Semantic Scoring

Scoring in ASI is built upon four fundamental principles:

Principle	Description
Semantic Foundation	Evaluation is based on extracted meaning and the implementation mechanism, not on surface-level text features
Multi-Criteria	Each idea is assessed across independent criteria reflecting different dimensions of its value
Weighted Aggregation	Criteria carry configurable weights, enabling adaptation to specific organizational objectives
Unified Scale and Comparability	All ideas are mapped to a common metric, enabling objective comparison regardless of domain, language, or author

10.2. Core Evaluation Criteria

The central semantic AI assesses each idea across the following key criteria:

Criterion	Description
Novelty	Degree of differentiation from existing solutions in global memory
Uniqueness of Mechanism	The primary criterion—how the implementation mechanism differs from known approaches

Feasibility	Technical and organizational viability
Implementation Practicality	Assessment of complexity, cost, risks, and required resources
Semantic Clarity	Coherence and structural quality of the idea
Impact Potential	Potential influence on industry, organization, or technological landscape
Reinforcement Value	Capacity of the idea to enhance or complement other solutions
Evolution Potential	Opportunities for further improvement and scaling

10.3. Multi-Criteria Model and Weighting

Each criterion is scored on a unified scale (e.g., 0 to 5). A weighting system is then applied, determining each criterion's contribution to the final score.

Example baseline weight distribution (configurable per organization):

Criterion	Weight Range
Uniqueness of Mechanism	25–30%
Novelty	20%
Feasibility + Implementation Practicality	20–25%
Impact Potential	15%

Semantic Clarity	5–10%
Reinforcement & Evolution Potential	5–10%

This weighting scheme may be adapted to organizational strategic priorities—for example, increasing Feasibility weight for conservative companies or Impact weight for innovation-focused divisions.

10.4. Final Score Formation Process

Scoring formation proceeds through multiple stages:

1. Semantic Structure Extraction — AI identifies key components of the idea
2. Independent Criterion Assessment — Each criterion is evaluated separately
3. Weight Application — Weighted sum of criteria
4. Final Score and Recommendations Generation — Consolidated output
5. Status Assignment — Classification such as "High Priority," "Synthesis Candidate," "Below Threshold," etc.

10.5. The Role of Scoring in the System

Semantic scoring performs several critical functions:

- Filtering — Elimination of weak and duplicate ideas at early stages
- Prioritization — Determining the order of analyst review
- Transparency — Providing substantiated evaluations to authors and analysts
- Synthesis Support — Identifying ideas that may reinforce one another
- Adaptation — Enabling configuration to specific organizational needs without altering the underlying reasoning model

10.6. Balancing Global Standardization and Local Adaptation

While the baseline reasoning model remains unified across the entire system, organizations have the flexibility to configure:

- Criterion weights
- Threshold values
- Prioritization and shortlist rules

This approach ensures global comparability of ideas while preserving local relevance for each organization.

Chapter 11. Architecture of Dynamic Novelty, Uniqueness, and Temporal Evolution

Novelty and uniqueness are not static properties of an idea. They are dynamic characteristics that change over time under the influence of incoming ideas, technological progress, and the evolution of the system's semantic memory.

The architecture of dynamic novelty assessment is one of the most important mechanisms ensuring the long-term relevance and fairness of the entire AI Semantic Infrastructure.

11.1. The Dynamic Nature of Novelty

The system evaluates novelty across several interrelated dimensions:

Dimension	Description
Semantic Novelty	Differentiation of meaning and purpose from existing memory
Architectural Novelty	Differentiation of the implementation mechanism—the primary criterion
Contextual Novelty	Relevance of the idea in the current technological and market context
Temporal Novelty	How far the idea leads or lags behind general industry development

The primary criterion for novelty assessment: the degree to which the implementation mechanism differs from all previously known solutions in the system.

11.2. Temporal Impact on Idea Evaluation

The system incorporates the temporal factor at all evaluation stages. An idea that was highly innovative two years ago may have lost significant novelty today due to:

- Emergence of more effective mechanisms

- Widespread adoption of similar solutions
- Shifts in technological context—new tools, materials, paradigms
- Knowledge accumulation within the relevant cluster

Each idea receives a timestamp, and the system periodically reassesses its position within the semantic landscape.

11.3. Evolution of Uniqueness

Uniqueness is a dynamic quantity that may:

Trajectory	Description
Decrease	When similar implementation mechanisms emerge
Increase	When industry development diverges in another direction
Transform	When the idea becomes a foundation for new syntheses or clusters
Disappear	When the mechanism becomes fully obsolete or superseded

The system continuously tracks the uniqueness trajectory of each accepted idea and utilizes this information when evaluating new submissions.

11.4. Influence of Accepted Ideas on New Evaluations

Accepted ideas actively shape the system's semantic landscape:

- They expand the space of known mechanisms, automatically reducing novelty of similar new ideas
- They establish context for evaluating subsequent proposals
- They may serve as a foundation for reinforcement and synthesis

The system distinguishes evolution of a mechanism (improvement, cost reduction, simplification) from mere duplication. Meaningful improvements are treated as new ideas, even when the underlying principle is similar.

11.5. Innovation Aging Architecture

The system maintains mechanisms for natural idea aging:

Aging Type	Description
Technological Aging	Emergence of more effective alternatives
Semantic Aging	Loss of problem or context relevance
Architectural Aging	Obsolescence of the implementation mechanism

Aged ideas may:

- Be deprioritized
- Be moved to archival or historical memory layers
- Be used primarily as context for evaluating new proposals

This enables the system to maintain the relevance and cleanliness of active semantic memory.

Chapter 12. Architecture of Semantic Tags, Indexing, Categorization, and Search

Semantic tags, indexing, and categorization form the structural layer of ASI that transforms the unordered flow of ideas into a navigable, coherent, and scalable knowledge base. Without effective organization of the semantic space, the system could not provide rapid search, relationship discovery, or synthesis support at scale.

12.1. Automated Semantic Tag Generation

The central semantic AI automatically generates semantic tags—structured markers that reflect the deep meaning of an idea, not merely surface-level keywords.

Primary tag types:

Tag Type	Description

Thematic	Application area and domain
Functional	Primary task and effect
Architectural	Implementation mechanism and key principles
Problem-Oriented	Problem being solved or constraint addressed
Contextual	Application conditions and requirements
Evolutionary	Development and synthesis potential

Tag generation principles:

- Fully automated, without user involvement
- Based on semantic structure, not lexicon
- Multi-layered and hierarchical
- May be updated over time as system knowledge evolves

12.2. Semantic Indexing

Semantic indexing transforms each idea into a multi-layered vector index, enabling the system to efficiently operate at scale with millions of records.

The index comprises:

Vector Type	Description
Semantic Vector	Meaning and purpose
Architectural Vector	Implementation mechanism
Contextual Vector	Conditions and constraints
Problem Vector	Task being addressed

Effect Vector	Expected outcome

Indexing occurs after central evaluation and is updated upon significant changes in global memory. This ensures high search speed and accuracy in comparing ideas.

12.3. Multi-Dimensional Categorization

ASI supports multi-dimensional categorization—each idea may simultaneously belong to multiple categories and subcategories.

Example categorization levels:

Level	Examples
Domain	Medicine, Energy, Transportation, Education
Functional	Automation, Optimization, Security
Architectural	Algorithms, Materials, Interfaces, Processes
Problem-Oriented	Cost Reduction, Reliability Enhancement, Scaling

Automated categorization enables ideas to exist across multiple contexts simultaneously—particularly valuable for interdisciplinary innovations.

12.4. Search Mechanisms

Search in ASI is semantic search—oriented toward meaning rather than literal keyword matching.

Primary search types:

- Search by meaning and purpose
- Search by implementation mechanism
- Search by problem
- Search by expected effect
- Search by semantic tags and categories
- Hybrid search (combination of multiple dimensions)

Results are ranked by semantic proximity, accounting for novelty, uniqueness, and temporal context. Search operates across both global memory and organizational nodes.

12.5. Knowledge Navigation

Navigation provides an intelligent interface for exploring the system's semantic space, including:

- Semantic clusters and relationship graphs
- Dynamic category maps
- Idea evolution pathways
- Synthetic nodes—points where multiple ideas converge
- Semantic landscape change history

Navigation is dynamic: as new ideas arrive, the knowledge structure automatically updates to reflect the current state of the innovation space.

Chapter 13. Architecture of AI Memory, Knowledge Storage, and Evolution

Memory is one of the most critical components of ASI. It transforms the stream of incoming ideas into a structured, continuously evolving knowledge base that serves as the foundation for novelty assessment, relationship discovery, and synthesis support.

ASI memory is not a static database, but a dynamic, self-updating semantic structure capable of compression, optimization, and evolution.

13.1. Memory of Processed Ideas

Each idea that successfully passes central evaluation is retained in global memory as a structured semantic object.

What is stored is not the original text, but a multi-layered semantic model comprising:

Component	Description
Semantic Vector	Meaning and purpose

Architectural Vector	Implementation mechanism
Contextual Vector	Conditions and constraints
Semantic Tags and Categories	Classifications and markers
Relationships	Links to other ideas
Timestamp and Change History	Temporal metadata and evolution log
Status and Evaluation Results	Outcome of assessment

This representation enables the system to efficiently compare new ideas against accumulated knowledge, determine novelty, and identify synthesis opportunities.

13.2. Storage of Rejected Idea Extracts

Rejected ideas do not disappear entirely. The system retains their semantic extract—a compact representation containing:

- Core meaning and implementation mechanism
- Reason for rejection
- Degree of similarity to existing ideas
- Timestamp

Purposes of extract retention:

- Prevention of duplicate resubmissions
- Improved understanding of idea evolution
- Enriched context for future evaluations
- Preservation of potentially useful fragments for possible synthesis

Personal author data and the full text of the idea are not retained.

13.3. Long-Term Knowledge Storage

Long-term memory is designed for retention of the most valuable elements:

- Accepted ideas and synthesized solutions
- Key architectural patterns
- Significant implementation mechanisms

- Evolving semantic clusters

Long-term storage principles:

- Storage in semantic, not textual form
- Multi-layered data organization
- Versioning and evolution support
- Efficient compression of aging records
- Persistent connectivity through the knowledge graph

13.4. Memory Optimization

To ensure scalability and performance, the system applies active memory optimization:

Optimization Technique	Description
Semantic Compression	Transformation of older ideas into key mechanisms and tags
Structural Merging	Consolidation of duplicate records into unified semantic nodes
Low-Value Content Pruning	Removal of fully obsolete or irrelevant structures
Knowledge Graph Restructuring	Reinforcement of important connections and weakening of weak ones
Dynamic Weight Redistribution	Increased significance of relevant ideas, decreased significance of outdated ones

13.5. Knowledge Base Evolution

The ASI knowledge base is a living, evolving system. It evolves under the influence of new ideas as follows:

Driver	Effect
Emergence of New Mechanisms	Creation of new nodes and relationships in the knowledge graph
Obsolescence of Mechanisms	Transfer to historical layer
Successful Syntheses	Generation of new, higher-level semantic entities
External Context Shifts	Reassessment and reorganization of existing knowledge in response to changes in technology, markets, or constraints

Through continuous evolution, the system's memory remains relevant and reflects the current state of the global innovation landscape.

Chapter 14. Architecture of Semantic Synthesis and Idea Amplification

Semantic synthesis and amplification are mechanisms that enable the system to extract additional value from weaker ideas without violating principles of fairness and transparency.

The system employs strict, predictable rules to avoid disputes over contribution, manipulation, and perceptions of injustice.

14.1. Philosophy of Synthesis and Amplification

Synthesis in ASI is not a tool for author convenience, but an architectural mechanism for improving the quality of the innovation signal.

Core principle: Synthesis is performed only if the resulting combined idea meets the established quality threshold. If synthesis does not meet the threshold, it is not created at all.

14.2. Amplification of Weak Ideas

Amplification occurs when a weak idea—one that does not independently meet the threshold—contributes a useful improvement to a strong idea.

Amplification rules:

Rule	Description
Weak idea receives no share in the primary idea	No attribution or ownership
Author of weak idea receives a fixed contribution of +1 to reputation	One-time per idea
Primary idea retains 100% contribution for its author	Full attribution preserved

This approach incentivizes participants to propose improvements while fully protecting the priority of the strong idea's author.

14.3. Semantic Combination (Synthesis)

Synthesis is the creation of a new idea through combination of two or more source ideas.

Mandatory condition: The combined idea must meet the quality threshold. If after synthesis the resulting idea does not meet the threshold, the synthesis is not performed and is not retained.

Synthesis rules:

- Mechanisms must be compatible and form a coherent, logical structure
- Synthesis is performed only if the result is significantly better than the source ideas individually

14.4. Contribution Accounting Model

ASI employs a strict and transparent contribution accounting model that minimizes disputes and prevents reputation inflation.

Combination Type	Conditions	Contribution Split	Reputation Impact
Weak + Weak	Synthesis possible only if combined idea meets threshold	50/50	Both authors receive +1 to rating
Strong + Strong	Synthesis created as additional option or improved version. Original authors retain full credit for source ideas	Full credit preserved for each	No additional rating for synthesis itself—contributions already accounted
Strong + Weak (Amplification)	Weak idea enhances strong idea (adds useful element, improves mechanism, reduces cost, etc.)	Strong author: 100%; Weak author: 0%	Strong author: +2 to rating; Weak author: +1 to rating

Important rule: Each author may receive positive contribution from a specific idea only once, regardless of how many syntheses or amplifications that idea subsequently participates in.

14.5. Synthesis Constraints

The system imposes strict constraints:

Constraint	Description
Quality Threshold	Synthesis performed only if resulting idea meets threshold
Architectural Compatibility	Combination permitted only with semantic and architectural compatibility
Meaningful Improvement	System does not create syntheses that yield no significant improvement
Idea Count Limit	Number of ideas participating in a single synthesis is bounded

These constraints protect system quality and prevent unnecessary combinations.

14.6. Architectural Significance of Synthesis

Synthesis and amplification enable maximum utilization of the global idea flow—including weaker ideas—while preserving the core principle: strong ideas retain full priority and author contribution.

The strict and transparent model ensures system behavior is predictable for all participants and minimizes conflicts regarding fairness.

PART IV. Evaluation, Review, and Human-AI Interaction

Chapter 15. Architecture of Evaluation and Idea Acceptance

The evaluation and idea acceptance architecture represents a strictly deterministic multi-layered pipeline that ensures objectivity, scalability, and system quality protection. It sequentially filters the idea stream, minimizing load on expensive components and routing only the most promising proposals to analysts.

15.1. Overall Pipeline Structure

Each idea passes through four primary evaluation levels:

Level	Component	Function
Level 1	Local Pre-Evaluation	Fast, lightweight filter executed by an isolated module. Tasks: filter obvious noise, spam, and low-quality content; provide early feedback to author
Level 2	Semantic Normalization	Central AI transforms free text into structured semantic representation—extracting purpose, implementation mechanism, context, constraints, and expected effect
Level 3	Deep Semantic Evaluation	Full analysis utilizing global memory: novelty determination, mechanism uniqueness, feasibility, synthesis potential, and multi-criteria scoring
Level 4	Final Decision and Routing	Based on scoring, current load, and organizational policies, the system determines idea status and subsequent path

15.2. Idea Progression Logic

The process is strictly sequential and includes the following key steps:

1. Pre-Evaluation — If the idea fails minimum requirements, the process terminates
2. Duplicate Check — If the implementation mechanism substantially overlaps with existing solutions, the idea is rejected
3. Novelty and Uniqueness Assessment — If the idea lacks sufficient novelty, it is rejected
4. Multi-Criteria Evaluation and Scoring — Overall score is generated
5. Synthesis and Amplification Check — System determines potential for combination or enhancement of other ideas
6. Final Decision with Load Consideration — System accounts for both idea quality and current organizational node capacity

15.3. Load Consideration and Dynamic Thresholds

One of the key architectural decisions is dynamic throughput management. Even good ideas may not make the shortlist under high load.

In such cases, the system:

- Classifies the idea as "Good, but did not pass due to capacity limits"
- Provides honest and useful feedback to the author, for example:
  "Your idea has been recognized as good, but did not make the shortlist today due to high load. We recommend resubmitting after some time."

This behavior is transparently explained to the user and motivates them not to abandon the idea but to try again when competition is lower.

15.4. Admission Criteria and Final Decision

For an idea to advance, it must satisfy minimum criteria: clear structure, feasible mechanism, semantic clarity, and absence of prohibited content.

The final decision is made exclusively by the central semantic AI. Possible statuses:

Status	Description
Rejected	Does not meet criteria or is a duplicate
Below Threshold	Has potential but insufficient strength
Passed Threshold	Directed to analyst queue

Synthesis/Amplification Candidate	System identifies potential for combination or enhancement

All decisions are accompanied by explanations accessible to the author (within the selected privacy mode) and to analysts.

Chapter 16. Architecture of Explainability and Decision Transparency

Explainability of decisions is one of the fundamental architectural principles of ASI. The system not only makes decisions but is also obligated to clearly explain why a particular decision was reached. This is essential for trust from authors, analysts, and organizations, as well as for auditability.

16.1. General Explainability Requirements

Every significant decision made by the central semantic AI must be:

Requirement	Description
Explainable	Contain human-understandable reasoning
Transparent	Show which factors and data influenced the outcome
Traceable	Enable reconstruction of the entire evaluation path
Documented	Retained for subsequent analysis and audit

Explanations are generated automatically and adapted to different user categories—author, analyst, auditor.

16.2. AI-Generated Explanations

The central AI generates structured explanations based on semantic analysis of the idea. Each explanation includes:

- Brief summary of the decision

- Primary reasons for the decision—strengths and weaknesses of the idea
- Assessment across key criteria with weighting rationale
- Comparison with the closest ideas in global memory
- Improvement recommendations—where applicable

Explanations are constructed exclusively from facts and semantic data, without subjective phrasing.

16.3. Transparency for Different Participants

The system provides varying levels of transparency based on user role:

Role	Access Level
Authors	Understandable explanations of decisions and recommendations for idea refinement
Analysts	Extended reasoning, including detailed mechanism breakdown, comparison with memory, and scoring rationale
Organizations and Auditors	Full traceability of the evaluation process and change history

Personal author data remains isolated in accordance with privacy principles (Chapter 6).

16.4. Evaluation Traceability

The system maintains a complete processing history for each idea, including:

- All evaluation pipeline stages
- Intermediate scoring results
- Status changes
- Participation in synthesis or amplification
- Timestamps and decision rationales

This history is retained for the period defined by data retention policies and may be restored at any time for audit or analysis.

16.5. Explanatory Note Generation

For analysts and organizations, the system automatically generates explanatory notes—structured documents containing:

- Semantic breakdown of the idea
- Detailed multi-criteria scoring
- Comparison with relevant ideas in memory
- Rationale for final status
- Recommended next actions

Notes are generated in a neutral, fact-oriented style and serve as the basis for human decision-making.

Chapter 17. Architecture of Analyst Workflows

While AI performs the primary filtering and initial evaluation, the final decision regarding the value and disposition of an idea remains with humans. Analysts represent the last and most critical human checkpoint in the ASI pipeline.

17.1. Role and Core Responsibilities of Analysts

The primary work of analysts encompasses three key areas:

Responsibility	Description
System Quality Control	Analysts verify the adequacy of AI selection and evaluation. They identify systematic filtering errors, false positives, and cases where strong ideas were missed or weak ideas were advanced
Final Expert Evaluation	Even if an idea passes all AI thresholds, the analyst determines whether it is genuinely needed by the organization at the current time, aligns with strategy, has practical value, and is implementable

Routing of Accepted Ideas	Analysts direct approved ideas to appropriate departments, product teams, or specific engineers capable of developing them

Analysts do not replace AI nor duplicate its work. They operate on top of AI results, serving as the final filter and distributor.

17.2. Analyst Workflows

A typical analyst workflow proceeds as follows:

1. Receive a prioritized queue of ideas with ready-made cards
2. Review the idea alongside AI-generated explanations
3. Verify system adequacy—initial assessment
4. Conduct own expert evaluation of value and relevance
5. Make a decision—Accept / Reject / Request Revision
6. Route accepted ideas to the appropriate department or team
7. Provide feedback to the author—when necessary

17.3. Review Queues and Prioritization

The system generates smart queues for analysts, prioritized according to:

Factor	Description
AI Overall Score	System-generated evaluation
Strategic Importance	Relevance to organizational priorities
Impact Potential	Potential influence on current projects
Urgency	Ideas related to critical issues—prioritized
Idea Age	Prevents stagnation and delays

Analysts may work with both general queues and thematic queues—by product, domain, or priority.

17.4. Analysis Interfaces and Analyst Workspace

The analyst workspace includes:

- Idea Card — Full text with brief AI explanation of why the idea passed the threshold
- Comparison View — Similar ideas from memory for reference
- Routing Tools — Quick selection of team or responsible individual
- Feedback Tools — Templates and free-form comments for authors
- Processing History — Complete audit trail of the idea
- System Error Flagging — Ability to mark system errors for improvement

The interface is designed for high-speed decision-making while preserving analysis depth.

17.5. Decision Support

AI provides minimal support to the analyst:

- Brief rationale for why the idea passed the threshold
- Display of closest comparable ideas
- Suggested routing directions

All responsibility and final decision-making rests exclusively with the analyst. AI does not participate in final value assessment and cannot override human decisions.

Chapter 18. Architecture of Human-AI Interaction

Human-AI interaction in ASI is built upon the principle of strict role and responsibility separation. AI performs large-scale analytical work but never makes final decisions. Humans remain the sole subjects with the right of final choice and accountability for consequences.

18.1. Principles of Human Control

Human control is a fundamental architectural principle of the system. Key provisions:

- Final decisions on every idea are always made by humans
- AI has no authority to independently approve, reject, or route ideas for implementation
- Humans exercise oversight over AI correctness

- AI serves as a powerful analytical and data preparation tool, not as a decision-making authority

This separation protects the system from automation errors and preserves human accountability at critical stages.

18.2. Responsibility Distribution

AI is responsible for:

- Semantic analysis and deep evaluation of ideas
- Novelty, uniqueness determination, and scoring
- Duplicate detection and synthesis potential identification
- Explanation generation and material preparation for analysts

Humans are responsible for:

- Verification of AI correctness
- Final expert evaluation of ideas in organizational context
- Decision-making regarding idea disposition—accept, reject, request revision
- Routing accepted ideas to appropriate departments
- System oversight and behavioral correction

18.3. Collaborative Idea Evaluation

Human-AI interaction is organized as a sequential process:

Stage 1. AI Processing — AI performs full analysis, generates scoring, explanation, and recommendations. AI work concludes at this stage.

Stage 2. Analyst Processing — Analyst reviews the idea alongside AI-prepared materials, conducts contextual assessment, and renders final decision.

Stage 3. Routing — Upon positive decision, analyst directs the idea to the appropriate department or team.

AI does not participate in Stages 2 and 3 and cannot override human decisions.

18.4. Final Decisions and Oversight

All final decisions—acceptance, rejection, implementation routing—are made exclusively by humans. AI possesses no veto or adjustment authority over these decisions.

Oversight of AI performance is conducted by analysts on an ongoing basis. They document:

- False positives and false negatives
- Systematic filtering errors
- Cases of insufficient or excessive stringency

This data is used for AI model improvement, but the oversight process itself remains exclusively human.

Chapter 19. Architecture of Feedback and User Experience

Feedback is the sole communication channel between the user and the system after idea submission. The feedback architecture is intentionally minimalist: it provides the author with sufficient information to understand their idea's status, but does not create false expectations or reveal internal processes.

19.1. General Principles of Feedback

Feedback in ASI is built upon the following principles:

- Minimalism — Users receive only necessary information
- Automated Component Transparency — Clear explanation of AI decisions
- False Expectation Protection — The system does not promise what it cannot guarantee
- Privacy Preservation — Particularly in anonymous modes
- Automated-Human Separation — Users clearly understand where AI work ends and human work begins

19.2. User Notifications

Users receive only notifications related to automated processing:

- Acknowledgment of idea receipt
- Local pre-evaluation results
- Notification of threshold pass or fail by the central AI
- Notification that the idea has been transferred to analysts—if it passed the threshold

The system does not notify about:

- Analyst review timelines
- Internal queue status
- Analyst or departmental decisions

19.3. Review Results

Users see only the results of automated AI evaluation:

- Passed Threshold — Idea transferred to analysts
- Failed Threshold — With brief explanation of reasons

If the idea passed the threshold, the author additionally receives:

- Brief semantic explanation from AI
- Rating change projection

Analyst decisions and their reasoning are not automatically disclosed to users. This protects analysts and maintains the system's focus on objective automated evaluation.

19.4. Rating Information

Rating is the primary long-term signal for the user. They see:

- Current reputation level
- Rating changes following threshold passage or synthesis participation
- Brief contribution history

Rating updates occur only upon threshold passage or successful amplification/synthesis. Users do not see internal calculations or contributions of other participants.

19.5. Author Communication

Communication with authors is possible but not guaranteed and is not mandatory. Analysts may contact the author only in the following cases:

- To clarify idea details
- When the idea has high value and additional information is needed
- To request revisions

The system explicitly communicates to users that analyst feedback is a bonus, not a guaranteed service. This approach prevents disappointment and maintains realistic expectations.

Chapter 20. Architecture of the Reputation System

The reputation system is a fully automated, transparent, and objective mechanism that reflects the real quality of a user's contribution to the innovation process. Reputation does not depend on the number of ideas submitted, author popularity, or human opinion. It is a pure function of idea quality, determined exclusively by AI and analyst results according to strict, pre-defined rules.

20.1. Reputation System Principles

- Reputation is fully automated and cannot be manually altered
- It does not depend on user activity or number of submitted ideas
- Human factors—analyst opinions about the individual—are excluded
- Each idea may yield only one final evaluation for its author
- The final evaluation is always equal to the maximum possible

20.2. Reputation Award Mechanics

Reputation is awarded only after an idea passes through central evaluation. The following scores are possible:

- +1 — A weak idea that passed the threshold through synthesis, or a weak idea that amplified a strong idea
- +2 — A strong idea that independently passed the threshold without synthesis. This is the baseline positive score for a quality, unique idea
- +3 — The maximum score, awarded by analysts. Analysts assign it anonymously, without seeing the author. After author disclosure, the score cannot be changed

Important rule: If an idea receives a +3 from analysts, all prior automated scores (+1 or +2) are overwritten. The final score is always the maximum.

20.3. Positive Scores and Their Meaning

- +1 — A weak idea that passed the threshold only through synthesis, or a weak idea that amplified a strong idea. The author receives this score once per idea
- +2 — A strong idea that independently passed the threshold without synthesis. This is the baseline positive score for a quality, unique idea
- +3 — The maximum score an idea can receive from analysts. Analysts assign it anonymously, without seeing the author. After author disclosure, the score cannot be changed

20.4. Negative Scores

The only way to receive a negative score in the system is -1 for ignoring a serious local pre-evaluation warning—submitting obvious noise, spam, or prohibited content despite the warning.

Analysts do not have the authority to assign negative scores. Negative scores are issued exclusively automatically.

20.5. Synthesis Contribution Accounting

- Weak + Weak — Upon successful synthesis that passes the threshold, both authors receive +1
- Weak + Strong — The weak idea receives +1, the strong idea retains its +2
- Strong + Strong — Synthesis is treated as an additional option; original authors retain their scores

Each author may receive positive contribution from a specific idea only once.

20.6. Rating Structure and Contribution Quality

The system stores not only the total rating sum but also the contribution structure:

- Number of ideas with +1 score
- Number of ideas with +2 score
- Number of ideas with +3 score

This enables visibility not just into the total sum but into the quality of the user's contribution: one author may have a high rating from many weak ideas, while another may have a high rating from fewer but very strong ideas.

This transparency makes reputation an honest and informative signal for organizations.

Chapter 21. Architecture of Abuse Protection and Adversarial Defense

ASI is a global infrastructure operating on a massive flow of ideas from millions of users, including corporations and government entities. The security architecture is therefore built upon principles of zero trust, multi-layered defense, and resilience against sophisticated attacks.

21.1. General Security Principles

The system is protected against the following threat types:

- Reputation inflation and manipulation
- Coordinated group attacks
- Semantic and technical manipulation
- Malicious and destructive submissions
- Attempts to undermine system integrity

Protection is implemented at all levels—from pre-evaluation to global memory and the analyst circuit.

21.2. Reputation Inflation Protection

Reputation is fully protected against manipulation through the following mechanisms:

- Each idea may yield only one final score
- The final score is always the maximum possible (+1, +2, or +3)
- Analyst scores (+3) overwrite prior automated scores
- Negative scores (-1) are issued automatically only for ignoring serious pre-evaluation warnings
- Analysts have no authority to manually modify ratings

Through these rules, reputation inflation via large numbers of weak ideas or syntheses is practically impossible.

21.3. Coordinated Attack Countermeasures

The system actively detects and neutralizes coordinated group actions:

- Graph-based analysis of user behavior and ideas
- Detection of anomalous submission patterns—frequency, mechanism similarity, synchrony
- Automatic segregation of suspicious streams
- Enhanced filtering for accounts and IPs with suspicious activity

Upon detection of coordinated activity, the system may temporarily reduce the priority of ideas from the group or strengthen pre-evaluation requirements.

21.4. Manipulation Detection

The system is protected against various forms of manipulation:

- Semantic — Attempts to deceive the model through phrasing
- Behavioral — Mass repetitive patterns
- Technical — Attempts to bypass filters

The primary defense mechanism is analysis of implementation mechanisms, not text. Even if the text differs substantially, the system recognizes recurring architectural logic and blocks pollution attempts.

21.5. Handling Malicious Submissions

Malicious, destructive, or prohibited ideas are blocked at early stages:

- Pre-evaluation detects and blocks obvious harmful content
- Central AI performs additional screening for dangerous mechanisms
- Malicious ideas do not enter synthesis, do not affect memory, and do not award reputation

Upon detection of systematic attempts to submit malicious content, the system applies strict restrictions to the source.

21.6. System Integrity Protection

System integrity is ensured at the architectural level:

- No manual modification of ratings, thresholds, or synthesis
- Strict circuit separation—identification, reputation, semantics, analytics
- Author anonymity from analysts until final decision
- Zero trust for all incoming data
- Continuous anomaly monitoring and automated response

Even in the event of compromise of individual components, the system preserves the integrity of core mechanisms—evaluation, reputation, and memory.

Chapter 22. Architecture of Governance, Organizational Flexibility, and Corporate Configuration

One of the key requirements for ASI was the ability to operate equally effectively in a small startup, a large corporation, or a government agency. To achieve this, the governance architecture is built upon the principle of a rigid, immutable core and high flexibility at the organizational process level.

22.1. Core Governance Principles

The ASI governance architecture rests on a clear separation:

- Global Core (Immutable) — Evaluation logic, synthesis, novelty, reputation, and memory
- Organizational Layer (Configurable) — Policies, routing, priorities, and processes specific to each organization

Organizations may flexibly configure the system to their needs but cannot modify the fundamental AI mechanisms. This ensures a unified quality standard and evaluation fairness across the entire global infrastructure.

22.2. Organizational Policies

Each organization defines its own policies, which affect:

- Priority areas and themes
- Internal idea routing rules
- Analyst access levels and roles
- Strategic filters and prioritization weights
- Author communication policies

Organizations are not permitted to modify:

- AI evaluation criteria and logic
- Synthesis and amplification mechanics
- Reputation award rules
- Baseline model thresholds

22.3. Configurable Evaluation Parameters

Organizations may adapt only surface-level parameters that do not affect the core:

- Weights of additional organizational criteria
- Shortlist prioritization rules
- Internal tags and categories
- Departmental auto-routing thresholds

The baseline multi-criteria evaluation model, key criterion weights, and central AI logic remain unchanged for all system participants.

22.4. Governance Models

The system supports several governance models, which organizations select independently:

Model	Description
Centralized	All ideas pass through a single analyst center—suitable for small companies and startups
Decentralized	Each department has its own analyst group—suitable for large corporations
Hybrid	Combination of centralized and decentralized models
External Experts	Engagement of external analysts or consultants—commonly used in government and international projects

Governance model selection does not affect semantic core AI operation.

22.5. Operational Flexibility

The organizational layer provides high operational flexibility:

- Dynamic restructuring of departments and responsibility matrices
- Creation of temporary priority areas and programs
- Flexible idea routing configuration
- Scaling of analyst counts and queues
- Adaptation to changing strategic company objectives

Any organizational configuration changes do not affect the global semantic core and do not compromise the objectivity of baseline idea evaluation.

Chapter 23. Architecture of Audit, Transparency, Security, and Access Control

ASI is a global critical infrastructure operating on ideas from millions of users and affecting the interests of corporations and states. The architecture of audit, transparency, security, and access control is therefore built upon principles of zero trust, event immutability, and maximum traceability with strict confidentiality preservation.

23.1. Audit Logs

All significant system events are recorded in an immutable, cryptographically protected audit log.

Recorded events include:

Central AI actions—evaluation, scoring, synthesis, memory updates

Analyst decisions and idea routing

Organizational policy and configuration changes

Filter bypass attempts and anomalous activity

All idea status changes

Logs possess the following properties:

Append-only structure—records cannot be deleted or modified

Cryptographic hash chain for tamper detection

Distributed storage with replication

Access strictly restricted by role and requiring multi-factor authentication

23.2. Transparency Mechanisms

System transparency is implemented at multiple levels:

Authors — Understandable AI decision explanations and current idea status

Analysts — Complete AI reasoning, processing history, and comparison with similar ideas

Organizations — Aggregated system performance reports, filtering effectiveness, and incoming flow quality

Oversight Bodies — Access to anonymized audit logs and statistics

The system ensures transparency of operational logic but never discloses author personal data without explicit permission and strict control.

23.3. Security Principles

ASI security is built upon the following principles:

Zero Trust — Every action and request is verified

Least Privilege — Minimum necessary access

Strict Circuit Segmentation — Identification, reputation, semantics, analytics

Cryptographic Protection — All sensitive data and logs

Continuous Monitoring — Real-time anomaly and threat detection

Automated Incident Response

The system is designed with awareness of threats at the nation-state actor level, including supply chain attacks, insider threats, and AI model compromise attempts.

23.4. Access Control

Access control is implemented through a multi-level role-based model (RBAC) with dynamic checks:

Role Access Permissions User Submit ideas and view status of own submissions Analyst Access to ideas and AI explanations; does not see author identity prior to decision Organization Administrator Policy, queue, and routing management; cannot modify AI logic or global mechanisms System Administrator Infrastructure responsibility; no access to semantic data or ratings Auditor / Oversight Body Access only to audit logs and anonymized reports All actions are logged and tied to role, time, and context.

23.5. Accountability and Regulatory Compliance

Accountability within the system is clearly distributed:

AI — Responsible for objective evaluation and explanations

Analysts — Responsible for final decisions and routing

Organization — Responsible for internal processes and policy adherence

Infrastructure Operators — Responsible for technical security and reliability

The system is designed in accordance with international standards—ISO 27001, GDPR and equivalents—data protection requirements, and critical infrastructure regulatory norms.

Chapter 24. Architecture of Queues and High-Load Infrastructure

ASI is designed as a global platform capable of processing millions of ideas daily from millions of users and organizations. The stream processing and resource management architecture is built around principles of asynchronicity, prioritization, level isolation, and horizontal scalability.

24.1. Processing Large Idea Streams

The system is engineered for extreme loads and guarantees that no idea is lost even at peak values.

Core stream processing principles:

- Asynchronous Ingestion — Ideas are immediately recorded and assigned a unique identifier. Users do not wait for synchronous results
- Stream Separation — Incoming ideas are logically partitioned by type—public, corporate, anonymous, priority—but pass through a unified standardized pipeline
- Level Isolation — Pre-evaluation, deep evaluation, synthesis, and analyst queues operate on separate resource pools, ensuring overload on one level does not impact others

24.2. Queue Management

Queues are the primary tool for load control and processing fairness.

The system maintains several specialized queues:

- Pre-Evaluation Queue — Input buffer for all incoming ideas—fast, high-throughput layer
- Deep Evaluation Queue — Only ideas that pass pre-evaluation; central semantic AI operates here
- Analyst Queue — Ideas that pass the threshold; may be partitioned by organization, domain, and priority
- Synthesis Queue — Separate stream for generation and processing of synthesized ideas
- Low-Priority Queue — Non-urgent but not rejected ideas; processed when resources are available

All queues are managed automatically using dynamic priorities and backpressure mechanisms.

24.3. Resource Allocation and Dynamic Scaling

Resources are distributed across pipeline levels dynamically:

- Pool Separation — Pre-evaluation, deep evaluation, synthesis, and analyst queues use isolated compute resources
- Automatic Scaling — Under load increase, the system horizontally adds nodes to the most heavily loaded layer
- Guaranteed Resources — Critical streams—from strategic partners, for example—may have reserved capacity
- Elasticity — Under load decrease, resources are automatically released for cost efficiency

24.4. Computational Efficiency

To ensure cost-effectiveness and performance, the following approaches are applied:

- Central AI Work Minimization — Only ideas that pass pre-evaluation reach the expensive stage
- Semantic Caching — Recurring mechanisms and structures reuse previous analysis results
- Deferred Processing — Non-critical tasks—extended reports, deep re-indexing—are executed in background
- Model Optimization — Different-capacity AI models are used at different pipeline stages

24.5. Infrastructure Scalability

The architecture is designed from the outset as horizontally scalable and multi-site:

- Addition of new compute nodes without system downtime
- Geographic distribution capability—regional installations with unified logical core
- Resilience to sudden load spikes—mass campaigns, government initiatives
- Support for both cloud and on-premise deployments

Through these solutions, the system can grow in tandem with the volume of the global innovation flow.

Chapter 25. Architecture of ASI as a Global Semantic Infrastructure

25.1. ASI as a Global Infrastructure Architecture

ASI defines the architecture of a global semantic infrastructure that any operator may deploy as an independent ASI node.

Each ASI node represents a fully autonomous installation of the system, operating according to unified architectural principles of semantic idea processing. The node provides infrastructure for organizations and individuals to connect regardless of geographic location, industry, or scale.

The ASI architecture does not presuppose the existence of a single global operator. On the contrary, it accommodates multiple independent operators, each deploying its own global infrastructure.

Examples of such operators include:

- Google ASI
- Microsoft ASI
- OpenAI ASI
- Other organizations implementing the ASI architecture

All ASI nodes follow unified architectural principles but remain completely independent infrastructures.

25.2. Organization Onboarding

Organizations gain access to ASI infrastructure through the open API of their chosen ASI node.

Prior to operation, the organization undergoes a registration process in which its legal status is verified and authentication credentials—API keys or equivalent mechanisms—are issued.

Upon connection, the organization is able to:

- Receive ideas addressed directly to it
- Receive ideas from the common flow of its ASI node
- Interact with the semantic infrastructure in accordance with its architectural rules

Each organization connects to a specific ASI node and becomes a participant in its local ecosystem.

25.3. Idea Submission Modes

Within each ASI node, two submission modes are supported.

Directed Submission

An idea is sent to a specific organization and is accessible only to that organization.

This mode is used when the author has identified the intended recipient in advance.

Common Flow

An idea is published to the common flow of a specific ASI node.

All organizations connected to that node may review such ideas and independently decide on further engagement with the author.

The common flow is local to each ASI node.

Ideas published in one ASI node are not visible to organizations operating through other ASI nodes.

25.4. Individual Interaction with ASI

Individuals may register with a chosen ASI node and utilize both submission modes:

- Directed submission to a specific organization
- Publication to the common flow of the node

Upon submission, ASI automatically performs semantic processing, quality validation, novelty assessment, classification, and routing.

All interactions between authors and organizations occur through the ASI infrastructure. Direct communication between parties is not part of the system architecture.

Authors may use either verified or anonymous accounts. Anonymous ideas contain no author information and carry no compensation mechanism.

25.5. ASI Node Autonomy

Each ASI node operates as a completely independent infrastructure.

The ASI architecture does not provide for automatic exchange of:

- Ideas
- Semantic memory
- User reputation
- Organizational internal data
- Processing results

Each operator independently:

- Maintains semantic memory
- Manages connected organizations
- Supports its own reputation system
- Processes all incoming ideas

Autonomy is a fundamental architectural principle of ASI, ensuring operator independence, infrastructure sovereignty, and participant data protection.

25.6. Global Nature of the Infrastructure

Each ASI node is a global infrastructure.

This means that, absent legal or regulatory restrictions, any organization or individual may connect regardless of country, industry, or scale.

Global reach is achieved through:

- Open APIs for organizations
- Open registration for individuals
- Unified semantic processing principles
- Scalable architecture designed for global idea flow

Thus, each ASI node constitutes an independent global infrastructure.

25.7. Freedom of Operator Choice

The ASI architecture does not mandate the use of any single operator.

Each organization independently determines which ASI node to connect to.

Selection factors may include:

- Technological implementation characteristics
- Data processing requirements
- National and industry regulatory requirements
- Preferred ecosystem

The ability to choose among operators creates a competitive environment, prevents infrastructure monopolization, and encourages the development of diverse ASI implementations while maintaining unified architectural principles.

25.8. Boundary of the ASI Architecture

This book describes the architecture of an individual ASI node as an independent global infrastructure.

Interaction among multiple independent ASI nodes falls outside the ASI architecture and is treated as a separate infrastructure layer.

Facilitating such interaction requires a specialized protocol and a distinct architectural model, which are described in the book GSI (Global Semantic Infrastructure).

Thus, ASI defines the architecture of the individual node, while GSI defines the architecture of the network of independent ASI nodes and the rules governing their interaction.

Chapter 26. Architecture of Data Storage and Information Lifecycle Management

The ASI data storage architecture is built upon principles of strict segmentation, minimization of stored information, and clear separation between organizational data and system semantic memory. This ensures regulatory compliance, security, and long-term efficiency of the global innovation intelligence.

26.1. Storage Policies

Each organization independently defines its data storage policies within permitted boundaries. It may configure:

- Retention periods for user ideas and statuses
- Archiving and deletion rules for internal documents and routes
- Storage policies for analyst decisions and associated metadata

The organization is not permitted to modify:

- Audit log retention periods and conditions
- AI semantic memory storage policies
- System log and cryptographic event chain retention rules

Principle: The organization manages its own data. The system manages its semantic memory and critical logs.

26.2. Data Retention Periods

The system distinguishes several data categories:

Category	Retention Period
Operational Data — active ideas, queues, current statuses	Several months to 3 years, per organizational policy
Archival Data — completed processes, older ideas	3 to 10 years
System Data — audit logs, security logs, cryptographic chains	Indefinite
AI Semantic Memory	Permanent — serves as the foundation for the entire system

26.3. Archiving and Deletion

Archiving transfers inactive organizational data to cold storage with reduced availability and cost. Only organizational data is archived. AI semantic memory and audit logs are not archived.

Data deletion is permitted only upon expiration of retention periods, explicit organizational request, or regulatory requirements. Upon deletion:

- Original text and personal data are completely erased
- Semantic structures—vectors, mechanisms, relationships—may be retained in anonymized form to sustain AI operation

26.4. Information Lifecycle Management

The lifecycle of each idea comprises five primary stages:

1. Creation — Idea recorded in operational storage with unique ID assignment
2. Processing — Pre-evaluation, normalization, and deep evaluation
3. Completion — Final status assigned—accepted, rejected, routed, etc.
4. Archiving — Transition to cold storage after active phase completion
5. Deletion — Complete erasure upon retention expiry or request

The principle of minimal necessary storage applies at all stages.

26.5. System Semantic Memory Management

AI semantic memory is a separate, critically important system component that:

- Stores extracted semantic structures, implementation mechanisms, relationships, and patterns—not original text
- Contains no personal data or author original text
- Is not subject to deletion or archiving
- Is continuously optimized and evolves

Principle: Organizational data may be deleted according to organizational policies. System semantic memory persists, as it is the foundation of the collective innovation intelligence of ASI.

Chapter 27. Architecture of Long-Term Learning and System Evolution

ASI is not a static idea evaluation tool, but a self-evolving intellectual infrastructure. Its long-term evolution occurs without altering fundamental rules and principles, solely through knowledge accumulation, refinement of understanding, and optimization of internal processes.

27.1. Knowledge Accumulation

The system accumulates knowledge not as raw text, but in the form of structured semantic objects. It retains:

- Implementation mechanisms of ideas
- Semantic structures and patterns
- Relationships between concepts and solutions
- Contextual dependencies and semantic evolution

Principle: The system remembers meaning and mechanism, not original text or author. This enables compact, secure, and anonymized knowledge storage while preserving utility for future evaluations.

27.2. Intellectual Capability Development

The system's intelligence grows incrementally through:

- Expansion of semantic memory volume and diversity
- Accumulation of rare and complex implementation mechanisms
- Improved capacity to distinguish subtle differences between similar solutions
- Enhanced synthesis quality and hidden relationship discovery

With each processing cycle, the system develops a deeper understanding of what constitutes valuable innovation across domains and contexts.

27.3. Semantic Model Adaptation

Model adaptation occurs evolutionarily and under controlled conditions:

- Addition of new concepts and relationships to the knowledge graph
- Refinement and restructuring of existing semantic structures
- Gradual expansion of expertise domains
- Adjustment of internal weights and priorities based on accumulated experience

Changes are never disruptive—new knowledge is always layered upon existing knowledge, preserving continuity.

27.4. Long-Term Optimization

System optimization includes:

- Improved evaluation accuracy and reduced false positives
- Enhanced synthesis and amplification quality
- Computational efficiency and processing speed optimization
- Strengthened resilience against novel attack and manipulation types

Primary optimization sources include analyst feedback, system error analysis, and long-term trends in the global idea flow.

27.5. Growth of Innovation Intelligence

Innovation intelligence of ASI is the system's capacity to:

- Achieve deep understanding of idea mechanisms
- Identify non-obvious relationships and synthesis opportunities
- Assess long-term solution potential
- Gradually expand the boundaries of known innovation space

Over time, the system evolves from a simple filter into a global collective innovation intelligence with increasingly refined understanding of valuable technological ideas in the contemporary world.

Long-Term Evolution Principle: Fundamental rules and mechanisms remain unchanged. Only the depth of understanding and quality of system performance grow.

Chapter 28. Architecture of the Innovation Lifecycle and Semantic Knowledge Base

The innovation lifecycle in ASI is the complete path of an idea from submission to its long-term impact on the system's semantic knowledge base and organizational expertise. The architecture ensures the continuous transformation of disparate ideas into structured, evolving collective knowledge.

28.1. Idea Lifecycle

Each idea passes through a standardized lifecycle consisting of the following stages:

1. Submission and Pre-Evaluation — Idea recording, initial filtering, and metadata preparation
2. Deep Semantic Evaluation — Analysis of novelty, uniqueness, feasibility, and potential
3. Decision and Routing — Status assignment and direction to analysts or synthesis
4. Analyst Review — Final expert evaluation by human analysts
5. Completion and Recording — Acceptance, rejection, or inclusion in synthesis with semantic memory update
6. Long-Term Storage and Evolution — Transition to archive or active semantic knowledge base

This lifecycle ensures transparency, traceability, and continuous knowledge enrichment of the system.

28.2. Knowledge Accumulation

Knowledge accumulation occurs not through storage of raw text, but through extraction and preservation of semantic structures:

- Implementation mechanisms
- Architectural patterns
- Semantic relationships and clusters
- Contextual dependencies

Accepted and synthesized ideas enrich the global semantic memory. Rejected ideas are retained in compressed form (extracts) to prevent repetition and improve context for future evaluations. Thus, every interaction with the system increases its collective intelligence.

28.3. Innovation Tracking

The system maintains continuous innovation tracking at the level of:

- Individual ideas and their statuses
- Semantic clusters and directions
- Evolution of implementation mechanisms
- Novelty dynamics and technology obsolescence

Tracking enables analysts and organizations to see not only individual ideas but also trends, gaps, and growth points in the technological landscape. The knowledge graph automatically updates to reflect the current state of the innovation space.

28.4. Semantic Knowledge Base Formation

The semantic knowledge base is the heart of ASI's long-term value. It is formed as:

- A multi-layered graph of concepts and mechanisms
- A hierarchy of semantic clusters
- A network of relationships between ideas at different maturity levels
- A dynamic map of the innovation space

The knowledge base is continuously optimized: outdated structures are compressed, strong patterns are reinforced, and new mechanisms receive priority. It serves as the foundation for novelty assessment, synthesis generation, and strategic organizational analysis.

28.5. Organizational Expertise Development

One of the most significant outcomes of ASI operation is the development of organizational expertise. Through systematic interaction with the idea flow, organizations:

- Gain an objective view of external innovations
- Identify gaps and opportunities in their product roadmaps
- Accumulate corporate semantic memory
- Improve the quality of internal evaluation and decision-making processes

Over time, each organization using ASI develops its own more precise and expansive innovation intelligence, amplified by the system's global semantic base.

Chapter 29. Author Participation, Operators, and ASI Operating Principles

Economic and incentive models of ASI are not the central subject of this book. Only general principles and possible options are described here. Specific mechanisms for funding, compensation, licensing, and commercialization must be determined by the organizations and teams that will implement and maintain the system.
The core philosophy of ASI is voluntariness, openness, and freedom of choice. The system was designed primarily as infrastructure, not as a commercial product.

29.1. Principle of Voluntary Participation

Participation in ASI is entirely voluntary. The author independently determines, for each idea, the submission mode and the rights they are willing to grant to the system.

29.2. Author Participation Models

The ASI architecture supports several basic participation models. The primary and prioritized model is the open model, as ASI was created first and foremost as a free public infrastructure.

Open (Voluntary Contribution) — The primary model. The author voluntarily submits the idea to the common flow. The idea may participate in syntheses, amplify other ideas, and become part of the global semantic memory.

Protected (Patent) — A separate system module. In this mode, the full idea text never leaves the author's device and is not stored in ASI.
The system receives only the idea title, a brief description provided by the author, and the local pre-evaluation score.

ASI serves exclusively as a secure bridge between the author and interested organizations. Synthesis of ideas in the protected model is not possible.
Specialized Models — Including research collaboration, corporate closed installations, and government variants.

The author independently selects the participation model for each submitted idea.

29.3. Value for Authors

The primary value of participation in ASI is the opportunity to be heard. The system enables meaningful ideas to reach organizations and individuals capable of evaluating and implementing them, regardless of the author's status, country, or connections.
Additional values include:
Objective reputation based on idea quality
Opportunity to be noticed by research teams and companies
Participation in the development of the shared innovation space

29.4. Role of Operators

The ASI operator provides the technical infrastructure: semantic core operation, scalability, security, and system development. The operator may be a commercial organization, a non-profit structure, a consortium, a government body, or a university.

29.5. Diversity of Operating Models

The ASI architecture inherently accommodates different approaches to system operation—from fully open and non-commercial to corporate and hybrid models. This enables adaptation of ASI to different purposes and contexts.

29.6. Federative Future

In the long term, ASI may evolve as a federated network of independent operators maintaining compatibility of core principles. Such a model would enable a global innovation infrastructure without a single center of control.

29.7. Strategic Purpose of ASI

ASI is conceived as an infrastructural layer, analogous to the internet, but oriented toward innovation. Its primary purpose is to connect meaningful ideas with implementation opportunities and to accelerate humanity's collective technological progress.

Conclusion. The Future of Global Innovation Intelligence

ASI is not another idea platform. It is an attempt to solve a fundamental problem of our time—the structural fragmentation of innovation.

30.1. The Problem ASI Solves

Humanity produces an enormous volume of technological ideas, yet possesses almost no infrastructure capable of effectively collecting, evaluating, connecting, and advancing them to implementation. The best ideas often remain unheard not because of low quality, but because of the absence of a suitable channel and mechanism.

Between the birth of an idea and its realization lies a vast gap—a gap of access, expertise, resources, and attention. ASI was created to bridge this gap.

30.2. Innovation as a Global Resource

In the 21st century, ideas become a strategic resource. Unlike oil or capital, they are distributed globally, emerging in the most diverse conditions, and independent of the author's formal status.

ASI transforms this distributed resource into a structured, accessible, and equitable flow, where value is determined not by the origin of the idea but by its semantic content and potential.

30.3. Democratization of Innovation

The system creates an environment in which:

- Any idea can be objectively evaluated
- Evaluation depends on meaning and mechanism, not the author's name or status

- Strong solutions have a chance to be noticed regardless of who proposed them

This is not equalization of talent, but removal of artificial barriers in the path of strong ideas.

30.4. From Information Networks to Innovation Networks

The internet connected people and information.

ASI is the next step: connecting ideas with implementation opportunities.

The system moves from data transmission to innovation signal transmission—from information search to effective solution discovery and synthesis.

30.5. The Role of AI in Collective Intelligence

In ASI, artificial intelligence functions not as an inventor, but as a scalable intermediary:

- Filters and structures the vast idea flow
- Identifies novelty and connections
- Amplifies weak signals
- Assists humans in making informed decisions

AI does not replace human intelligence—it unlocks its potential at a global scale.

30.6. Long-Term Vision

ASI is not a single platform, but a new class of infrastructure. In the future, we may see:

- Multiple operators and federated networks
- Distributed semantic cores
- Global innovation ecosystems
- Systems operating for decades, evolving alongside the technological landscape

This is infrastructure capable of supporting humanity's collective innovation intelligence at a new level.

30.7. Final Word

The greatest ideas already exist—they are born every day in every corner of the world.

The primary challenge of our time is to build infrastructure that allows these ideas to be heard and realized.

AI Semantic Infrastructure is an architectural proposal for such infrastructure.

Not a final answer, but the first serious step toward a world where every strong idea has a real chance to shape the future.

Afterword

This book has described the architecture of AI Semantic Infrastructure (ASI)—a system designed to ensure that meaningful technological ideas are heard and given a chance at realization.

We have deliberately limited the scope of this work to one specific objective: the creation of an efficient, scalable, and fair infrastructure for idea processing. ASI was conceived as a practical and implementable system that can be deployed today—whether within a single organization or on a broader scale.

However, ASI represents only the first level of a larger process.

As multiple such systems emerge—across different companies, research centers, government organizations, and independent operators—the natural next step will be their compatibility and the gradual formation of a broader network.

In the future, there may not be a single ASI, but an entire ecosystem of semantic infrastructures operating under common principles while maintaining their independence. Such an environment would enable a truly global innovation space, in which ideas can find their way to those capable of realizing them, regardless of organizational and geographic boundaries.

We do not claim that the architecture described in this book is final. On the contrary, it is offered as a proposal—open to discussion, critique, and further development. Any infrastructure of this scale will inevitably evolve alongside technology, society, and our understanding of collective progress.

Perhaps future generations will perceive innovation search and coordination infrastructures as naturally as we today perceive the internet, email, or search engines.

If that happens, then the purpose of this book will have been achieved.

What we believe in remains unchanged:

Humanity deserves a system in which every truly valuable idea has a chance to be heard.

Appendices

Appendix A. Idea Evaluation Criteria System

ASI employs a multi-criteria semantic evaluation model that analyzes each idea across several independent dimensions. Each criterion is scored separately, after which a weighted final score is generated.
Core Criteria:
Novelty — Degree of differentiation from existing solutions in global memory

Mechanism Uniqueness — The primary criterion; how the implementation mechanism differs from known approaches

Feasibility — Technical and organizational viability of the idea

Practical Value — Potential benefit to the organization or industry

Scalability — Capability for expansion to other contexts

Potential Impact — Long-term effect on technology, economy, or society

Semantic Clarity — Logical coherence and comprehensibility of the idea

Risk and Complexity — Probability of failure and required resources

Criterion weights are configurable to the objectives of the specific organization, but the underlying reasoning model remains unified.

Appendix B. Idea Evaluation Examples

Example 1. New Battery Type

The AI identified an original energy storage mechanism. Novelty and mechanism uniqueness received high scores. The idea passed the threshold and was routed to analysts.

Example 2. Drone-Based Drug Delivery

The system detected overlap with existing logistics solutions. After mechanism analysis, the idea was recognized as an amplification of an existing cluster and routed to synthesis.

Example 3. New AI Optimization Algorithm

High mechanism novelty, strong feasibility, and significant impact potential. The idea received a high score and was placed in the priority queue.

Example 4. Idea Without a Clear Mechanism

Despite an interesting objective, the absence of a well-developed implementation mechanism led to rejection at the deep evaluation stage.

Appendix C. Examples of Semantic Synthesis

Source Ideas:

- Idea A: A weak idea for cargo delivery using small drones
- Idea B: A weak idea for a network of window-mounted charging stations

Synthesis Result:

New Idea: "Autonomous cargo delivery network using drones and

distributed window-mounted charging stations."

The system:

- Identified compatible mechanisms
- Formed a coherent architecture
- Verified that the synthesis passes the threshold
- Created a new semantic entity

Both authors received +1 to their reputation.

Appendix D. Reputation System Mechanics

Reputation is formed exclusively automatically based on evaluation results:

- +1 — A weak idea passed the threshold through synthesis or amplified another idea
- +2 — A strong idea independently passed the threshold
- +3 — High analyst evaluation—overwrites previous scores

Each idea yields only one final score for the author. The system stores the rating structure—count of +1, +2, +3—enabling assessment not only of the total sum but also of contribution quality.

Appendix E. Semantic Tags, Indexing, and Search

The system automatically generates multi-layered semantic tags and vector indices for each idea. Search operates by meaning and implementation mechanism, not by keywords. This enables efficient discovery of similar ideas, cluster construction, and dynamic navigation across the global knowledge base.

Appendix F. ASI Use Cases

- Independent Inventor — Idea validation, reputation growth, partner discovery
- University / Research Group — Hypothesis evaluation, collaboration discovery
- Startup — Idea validation, investor discovery, solution synthesis
- Corporation — External R&D, early access to innovations, patenting
- Government Agency — Strategic technology and solution monitoring at the national level

---

Appendix G. Typical ASI Configurations

- Corporate ASI — Closed system within a single company
- Industry ASI — Consortium of multiple companies within one industry
- Government ASI — For strategic national projects
- Global Open ASI — Public network for all authors

Appendix H. Comparison with Existing Systems

ASI is fundamentally different from:

- Corporate Idea Banks — Limited to internal employees
- Patent Systems — Operate only on formalized inventions
- Crowdsourcing Platforms — Lack deep semantic evaluation
- Accelerators — Focus exclusively on business models

The key differentiators of ASI are global scale, semantic mechanism evaluation, synthesis support, and dynamic memory.

Appendix I. ASI Risks and Limitations

Key risks and mitigation measures:

Risk	Mitigation
AI Errors	Minimized through multi-layer verification and human oversight
False Positives / False Negatives	Tracked by analysts; data used for model improvement
Centralization	Mitigated through support for distributed and federated models

Legal Conflicts	Clear separation of free and patent models; transparent rules
Manipulation	Prevented through graph analysis, semantic verification, and strict reputation rules

ASI was designed with awareness of these risks and contains built-in mechanisms for their minimization.